A Framework for Comparing Medical Knowledge Representation Systems

Case study: the representation of anatomy in the Foundational Model of Anatomy (FMA) and GALEN Common Reference Model (GALEN)

Final Report (04/08/2002 – 04/04/2005)

Songmao Zhang, Ph.D.
Mentor: Olivier Bodenreider, M.D., Ph.D.

Overview:

The general framework of this study is the comparison and evaluation of the structures and reasoning potentials in different medical knowledge representation systems. Two representations of anatomy were selected as a case study: the Foundational Model of Anatomy (FMA) and the GALEN Common Reference Model (GALEN), both comprehensive enough to support clinical applications. The FMA is modeled in a frame-based structure while GALEN is based on description logic. Our approach consisted of aligning concepts across systems based on their lexical resemblance and structural similarity of taxonomical and partitive relationships. Moreover, complex structural aligning rules were developed for identifying mappings between groups of concepts across systems. Additionally, associative relationships between the FMA and GALEN were compared based on the concept matches. In order to facilitate the alignment, knowledge implicitly embedded in concept names and combinations of hierarchical relationships was investigated in each system in addition to knowledge explicitly represented.

To evaluate our alignment approach, we compared it to a schema matching system developed at Microsoft for aligning the FMA and GALEN. Moreover, we applied our approach to aligning two other representations of anatomy: the Adult Mouse Anatomical Dictionary (MA) and the anatomy subset of the NCI Thesaurus (NCI). The resulting matches were reviewed manually by a domain expert.

We also studied the alignment of multiple ontologies by comparing pairwise direct alignment with alignment through a reference ontology. The FMA was selected as the reference and MA and NCI were aligned to the FMA directly by our aligning approach. An indirect alignment between MA and NCI was derived accordingly, and was subsequently compared with matches in the direct alignment between the two systems.

We proposed ontology modeling principles and investigated the degree to which the FMA complied with these principles. Furthermore, in order to compare frame-based structures versus description logics in supporting large-scale biomedical ontologies, we converted the FMA from its frame-based representation to the description logic language OWL DL. An OWL reasoner was used to identify inconsistencies in the FMA by reasoning on the necessary and sufficient conditions defined for classes as well as other properties. Additionally, from a different perspective we investigated the FMA by converting it into the Region Connection Calculus (RCC model). Using the RCC composition table, spatial relationships among concepts were determined or refined, and inconsistencies were detected in the partitive and spatial knowledge representation in the FMA.

Aligning the FMA and GALEN using lexical and structural methods [1,2]

❑ Methods of aligning FMA and GALEN

Aligning FMA and GALEN consists of the following four steps: 1) acquiring terms, 2) identifying anchors (i.e., shared concepts) lexically, 3) acquiring semantic relations, and 4) identifying anchors structurally. The first two steps constitute the lexical component of our method, and the last two the structural component.

o Lexical alignment

Lexical alignment compares the two systems at the term level, by exact match and after normalization. This process makes the source and target terms potentially compatible by eliminating such inessential differences as inflection, case, hyphen, and word-order variation. Both preferred terms and synonyms, when available, are used in the alignment process. Concepts exhibiting similarity at the lexical level across systems are called anchors, as they are going to be used as reference concepts in the structural alignment. For example, the concepts *Fibularis tertius* (synonym: *Peroneus tertius*) in FMA and *Peroneus Tertius* in GALEN were identified as anchor concepts. Moreover, UMLS synonymy was used to identify additional anchors. For example, *Greater pelvis* in FMA and *Major pelvis* in GALEN, although

lexically different, were identified as an anchor because they name the same anatomical concept in UMLS.

o Structural alignment: acquiring inter-concept hierarchical relations
Inter-concept hierarchical relations explicitly represented in each system were extracted. In addition, three techniques were applied for acquiring implicit knowledge: complementing, augmenting and inferring.

- Complementing

Hierarchical relationships IS-A and PART-OF have inverse relationships, INVERSE-IS-A and HAS-PART. Some relations are represented bi-directionally and some are not. For example, *<Arm, HAS-PART, Humerus>* is explicitly represented in FMA but its inverse relation is missing. In order to perform simple comparisons of semantic relations across systems, we complemented FMA and GALEN with the inverse relations that were not explicitly represented. For example, we generated the relation *<Humerus, PART-OF, Arm>.*

- Augmenting
Instead of representing a PART-OF relationship between concepts *X* and *Y*, some knowledge representation systems reify the PART-OF relationships in a concept named *Subdivision of Y* and assert an IS-A relationship between concepts *X* and *Subdivision of Y*. These two representations are equivalent for most purposes. However, nearly 98% of such reified PART-OF relationships are not explicitly represented in FMA [August 30, 2002 version]. *Cardiac chamber IS-A Subdivision of heart*, but there are no explicit hierarchical relationships between *Cardiac chamber* and *Heart* in FMA. We made such reified relations explicit. The augmentation was also applied to concepts such as *Organ component of X* in FMA. In order to facilitate aligning, similarly in GALEN, the augmentation was applied to concepts named including *SubdivisionOfX*, *ComponentOfX*, and *StructuralComponentOfX*.
Relations are also captured in various other linguistic phenomena such as nominal modification and prepositional attachment. The former often represents a hyponymic relation involving the head of the noun phrase. For example, a *Sweat gland* is a kind of *Gland*. Therefore, a *<X Y, IS-A, Y>* relation can be tentatively extracted from the term *X Y*. In anatomical terms, prepositional attachment using "of" (*X of Y*) often denotes a partitive relation between *X* and *Y*. For example, we generated the relation *<Upper lobe of lung, PART-OF, Lung>* from the term *Upper lobe of lung*.

- Inferring
Inference consists of generating new inter-concept relationships by applying inference rules. These inference rules combine IS-A and PART-OF relationships among concepts to generate additional PART-OF relations. They represent limited reasoning along the PART-OF hierarchy, generating a partitive relation between a specialized part and the whole or between a part and a more generic whole. For example, we add the relation *<First tarsometatarsal joint, PART-OF, Foot>* from *<First tarsometatarsal joint, IS-A, Joint of foot>* and *<Joint of foot, PART-OF, Foot>.*in FMA.

o Structural alignment: identifying similar and conflicting hierarchical relations among anchors across systems
With the explicit and implicit relations, structural alignment identifies structural similarity and conflicts among anchors. Structural similarity is defined by the presence of common relations among anchors across systems. In other words, for a given anchor, the relations to other anchors observed in FMA are expected to be present in GALEN (and vice versa). For example, *Fibularis tertius* in FMA and *Peroneus Tertius* in GALEN share hierarchical relationships to other anchors, including IS-A *Muscle of leg*, PART-OF *Leg*, and HAS-PART *Muscle cell*. Therefore, the lexical mapping between *Fibularis tertius* in FMA and *Peroneus Tertius* in GALEN is supported by structural similarity.
Conflicts are defined by semantic incompatibility between anchors. The first type of semantic incompatibility is represented by the presence of opposite relationships (e.g., PART-OF and HAS-PART) between two anchors across systems. The second type of semantic incompatibility is based on the disjointedness of top-level categories in knowledge representation systems. In practice, if anchor concepts have relationships to disjoint top-level categories in FMA and GALEN, they are declared semantically distinct and the corresponding anchor is removed. For example, *Foot* in FMA and *feet* in GALEN match after normalization, resulting in an anchor. However, the two concepts are considered semantically incompatible because FMA's *Foot* is linked to the top-level concept *Anatomical structure*, while GALEN's *feet* is linked to *Unit*, *Anatomical structure* and *Unit* being disjoint top-level categories.

❑ Results of aligning FMA [August 30, 2002 version] and GALEN [v.5]
2,353 pairs of matching concepts across systems were identified as anchors by lexical comparison. In structural comparison without implicit knowledge, 69% of the anchors were supported by positive evidence, 30% didn't have any structural evidence and 1% received negative evidence. Taking advantage of implicit knowledge increased the number of anchors acquiring positive evidence to 91%, decreased the number of anchors receiving no evidence to 7.5%, and also revealed a few more conflicts between anchors.

❑ Implementation
We have developed a unifying Java-based interface to access both FMA and GALEN data. The Java API provided by Protégé-2000 is used to communicate with FMA. The GALEN server is accessed through a COM interface and we use JACOB, a Java-COM bridge, to interact with GALEN in Java.

❑ Problems detected during the alignment
Some of the inconsistencies and errors in the two systems or their development tools were detected automatically during the alignment. They have been reported to the GALEN and FMA developers, respectively. Some of them are listed as follows.

   o Reflexive and circular hierarchical cycles, either by *IS-A* or *PART-OF*. For example in FMA, a *PART-OF* cycle was identified in *Apex of urinary bladder -PART-OF→ Urinary bladder - PART-OF→ Apex of urinary bladder.*

   o Combined hierarchical cycles, e.g., *X IS-A Y* and *Y PART-OF X*. For example in GALEN, a combined cycle was identified in *Conduction Fibres -IS-A→ Myocardium -PART-OF→ Sino Atrial Node -PART-OF→ Conduction Fibres Of Right Heart Atrium -IS-A→ Conduction Fibres.*

   o Implicit errors are detected in the implicit knowledge of the system. For example, *Body of nail IS-A Subdivision of body of nail* in FMA; *WristJoing IS-A UlnoCarpalJoint* which *IS-A ComponentOfWristJoint*; again in GALEN, and *MalePerineum IS-A ComponentOfMalePerineum* in GALEN.

   o Invalid semantic relations, i.e., relation *<X, relationship, Y>* is present in the system but concept *Y* does not exist in the system.

   o The COM of GALEN returns incomplete concept names or definitions when their lengths go beyond some limitation, which makes defective the functions of getting the concept by name and getting the name of concept.

Aligning the FMA and GALEN solely based on structural similarity [10]

❑ Structurally aligning concepts of FMA and GALEN that do not share lexical similarity
Based on the anchors supported by both lexical and structural similarity, techniques for acquiring new mappings based solely on structural similarity are being developed. In practice, the unmapped concepts (non-anchors) in FMA and GALEN were compared according to their structural connections to the anchors. We developed various structural aligning rules and applied them to identify one-to-one match, one-to-group match, group-to-group match, and concepts without match in another system.

   o Aligning Rule 1: Some high-level concepts in GALEN represent non-anatomical categories, e.g., *Process*, *NonNormalPhenomenon*, *RiskFactor*, *Graft*, *Food*, etc. For any concept in GALEN, if is a descendant of one of the non-anatomical categories, it is considered as non-anatomical. Such non-anatomical concepts in GALEN do not likely have any matches in FMA. For example, *SupernumeraryThumb* was labeled non-anatomical from being a descendant of *NonNormalPhenomenon*.

   o Aligning Rule 2: Two concepts across systems are likely to be a match if they reach the same nonempty descendant anchor set through *INVERSE IS-A*, directly or indirectly. For example, *Cuneiform* in GALEN and *Cuneiform bone* in FMA.

   o Aligning Rule 3: For any anchor *(X,Y)*, if *X* is a leaf concept in one system while *Y* is not in another system and all of Y's descendants *{Y₁, …, Yₙ}* are not anchors, then it is possible that Y_1, …, and Y_n do not have matches in the system of *X*. For example, the anchor concept *Ankle* in FMA has two children: *Left ankle* and *Right ankle*, while *Ankle* in GALEN is a leaf concept. *Left ankle* and *Right ankle* do not likely have matches in GALEN.

o Aligning Rule 4: For two concepts of an anchor *(X,Y)*, if *X* and *Y* reach the same children anchor set through direct INVERSE IS-A (possibly empty), and *X* and *Y* have the same number of non-anchor children, *{X₁, …, Xₙ}* and *{Y₁, …, Yₙ}*, respectively, then it is possible that either two groups of concepts match as a whole, i.e., *{X₁, …, Xₙ}* matches *{Y₁, …, Yₙ}*; or, the two groups of concepts match on a one-to-one base, i.e., *X₁* matches Y₁, …, *Xₙ* matches Yₙ.

- Some of them match as a whole group. For example, anchor *HeadOfRadius* in GALEN has two children: *DistalHeadOfRadius* and *ProximalHeadOfRadius*, while anchor *Head of radius* in FMA has two children: *Head of left radius* and *Head of right radius.* All these children concepts are non-anchors. Therefore, a group-to-group match was identified between *{HeadOfRadius, DistalHeadOfRadius, ProximalHeadOfRadius}* in GALEN and *{Head of radius, Head of left radius, Head of right radius}* in FMA. This shows the different views the two systems used in classifying head of radius.
- Some of them actually indicate a group of one-to-one matches. For example, *AtrialMyocardium* in GALEN and *Myocardium of atrium* in FMA composed an anchor. *AtrialMyocardium* has two children: *LeftAtrialMyocardium* and *RightAtrialMyocardium*, and *Myocardium of atrium* has two children: *Myocardium of left atrium* and *Myocardium of right atrium.* All these children concepts are non-anchors. Two one-to-one matches were identified in this case, i.e., *LeftAtrialMyocardium* matches *Myocardium of left atrium*, and *RightAtrialMyocardium* matches *Myocardium of right atrium*.

o Aligning Rule 5: For two concepts of an anchor *(X,Y)*, if *X* and *Y* have the same nonempty children anchor sets through direct INVERSE IS-A, all of *X*'s children are anchors, but *Y* has non-anchor children *{Y₁, …, Yₙ}*, then it is possible that concepts *{Y₁, …, Yₙ}* do not have any matches in the other system. For example, anchor *CerebralHemisphere* in GALEN has four children: *LeftCerebralHemisphere*, *RightCerebralHemisphere*, *NonDominantHemisphere* and *DominantHemisphere*, the first two being anchors. Anchor *Cerebral hemisphere* in FMA has two children: *Left cerebral hemisphere* and *Right cerebral hemisphere*, both being anchors matching the two anchor children under *CerebralHemisphere* in GALEN. Therefore, *{NonDominantHemisphere, DominantHemisphere}* in GALEN may not have any matches in FMA.

o Aligning Rule 6: If concept *X* is non-anchor in one system, all of *X*'s children are anchors *{X₁, …, Xₙ}*, X's closest anchor ancestor is *U*, in the other system, the corresponding anchor concepts matching *U*, *X₁*, …, *Xₙ* are *V, Y₁, …, Yₙ*, respectively, and V is the closest anchor ancestor *Y₁, …,*and *Yₙ* share, then it is possible that concept *X* matches a group of concepts *{Y₁, …, Yₙ}* in the other system. For example, anchor *Lobe of thyroid gland* in FMA has three children: *Right lobe of thyroid gland*, *Left lobe of thyroid gland* and *Pyramidal lobe of thyroid gland*, all being anchors. In GALEN, all these three anchors are descendants under anchor *LobeOfThyroidGland* as well, except that there is an additional intermediate concept *(Lobe which < isSpecificSolidDivisionOf ThyroidGland >)* that is a children of *LobeOfThyroidGland* and father of *LeftLobeOfThyroidGland* and *RightLobeOfThyroidGland*. According to Aligning Rule 6, *(Lobe which < isSpecificSolidDivisionOf ThyroidGland >)* in GALEN was matched to anchor group *{Left lobe of thyroid gland, Right lobe of thyroid gland}* in FMA.

o Aligning Rule 7: Suppose *mdes(X)* denotes the set of all anchors in the descendants of concept *X*. For any two concepts *X₁* and *X₂* in one system and concept *Y* in another system, if $mdes(X_1) \not\subseteq mdes(X_2)$, $mdes(X_2) \not\subseteq mdes(X_1)$, and $mdes(X_1) \cup mdes(X_2) = mdes(Y)$, then it is possible that the single concept *Y* matches a group of concepts *{X₁, X₂}*. For example, the concept *ExtremityLongPart* in GALEN has four anchors in the descendants: *Arm*, *Forearm*, *Leg* and *Thigh*. In FMA, *Proximal free limb segment* has two anchors in the descendants: *Arm* and *Thigh*, while *Middle free limb segment* has two anchors in the descendants: *Forearm* and *Leg*. A one-to-group match is identified between *ExtremityLongPart* in GALEN and {*Proximal free limb segment*, *Middle free limb segment*} in FMA.

o Associative relationships in GALEN were used to validate the identification of FMA concepts without matches. While all FMA concepts exhibit laterality (when appropriate) by names, e.g., using words "*left*" and "*right*" (e.g., *Ankle* has two children named *Left ankle* and *Right ankle*), many GALEN concepts are leaves with associative relationships specifying the property of laterality rather than being classified further (e.g., *Ankle* is a leaf with associative relationship *"isPairedOrUnpaired"* and the value is *"mirrorImaged"*, which is a kind of *"leftRightPaired"*). Such property of the GALEN concept *Ankle* validates that *Left ankle* or *Right ankle* in FMA can not be precisely mapped to any concept in GALEN.

- o Aligning Rule 8: If *X* is identified as having no matches in another system, then concepts named "*… of X*" or "*X of …*" may not have matches in another system as well. For *Left ankle* presented earlier, there are 28 concepts in FMA named "*… of left ankle*" including *Skin of left ankle*, *Anterior part of left ankle* and *Deep lymphatic vessel of left ankle*. These concepts are unlikely to have any matches in GALEN.

- o There is a large imbalance between FMA and GALEN in number of concepts that are not yet identified in alignment (either as matches or having no matches). The number is about 58,000 in FMA and 9,900 in GALEN. A lexical analysis was conducted to investigate how these concepts differ in names across systems. A word index was created from concept names in both systems, and words with large differences in frequency of occurrence across systems are being studied. For example, the word "proper" occurs in 1,667 FMA concepts but only once in GALEN. These FMA concepts represent parts of organs, arteries, bones, etc. (e.g., *Ureter proper*, *Hepatic artery proper*, *Subdivision of hepatic artery proper*, *Trapezium proper*). This may indicate that GALEN does not specify this specific entity in anatomy. Thus mappings can not be provided to these concepts in FMA.

- ❑ Results of aligning FMA [December 2, 2004] and GALEN [v.6] by lexical and structural methods
  So far, about 39% of FMA concepts and 84% of GALEN concepts were identified in the alignment. The alignment is composed of various types of matches, mismatches, and concepts without matches.

  - o One-to-one matches where two concepts across systems share lexical similarity as well as positive structural evidence (i.e., anchors). 3,097 matches.

  - o One-to-one conflicting concepts with negative structural evidence. 26 pairs.

  - o One-to-one matches eliminated after manually disambiguating the cases where two concepts in one system both match one concept in another system. 175 pairs.

  - o One-to-one matches where two concepts across systems share, although no lexical similarity, the same set of anchors in the descendants. 139 matches.

  - o One-to-group matches where two concepts in one system jointly have the same set of anchors in the descendants with one concept in another system. 32 matches.

  - o Concepts in GALEN that do not represent anatomical entities. 39,885 concepts.

  - o Concepts that do not have matches in another system due to the classification differences of anchor concepts in taxonomy. 9,302 concepts in FMA and 444 in GALEN. In 68% of the cases, the results can be validated by laterality terms (containing "*left*"/"*right*") in FMA and associative relationships in GALEN.

  - o Concepts that do not have matches in another system since lexically they represent parts or kinds of concepts without matches. 17,964 concepts in FMA and 62 in GALEN.

  - o One-to-group matches reflecting the different granularities the two systems used in classifying the same concept. 46 matches.

  - o Group-to-group matches either representing a group of one-to-one matches or reflecting the different views the two systems used in classifying the same concept. 41 matches.

Comparing the associative relationships between the FMA and GALEN [5]

- ❑ Comparing associative relationships
  There are 59 associative relationships in FMA and 562 in GALEN. Under our hypothesis, a correspondence between relationships across systems is indicated by the frequent association one relationship between two concepts in one system and either another relationship or a combination of relationships between the equivalent concepts in the other system. Thus, our method consists of identifying for each associative relationship in one system the relationship (or combination thereof) in another.

  - o Acquiring associative relations
    Acquiring associative relations consists of extracting the relations explicitly represented and complementing the missing inverse relations. For example, the explicit relation *<ArteryOfNasalPassage, serves, NasalCavity>* was extracted and its missing inverse relation *<NasalCavity, isServedBy, ArteryOfNasalPassage>* was complemented in GALEN.

  - o Augmenting associative relations
    Knowledge representation systems sometimes reify the associative relation *<X, branch_of, Y>* in a hierarchical relation between *X* and a concept called *Branch of Y*, i.e., *<X, isa, Branch of Y>*. These two

relations are semantically equivalent. In order to facilitate comparisons across system, in each system, we made explicit the relations implicitly embedded in the concept names (reified). We applied the augmentation technique to the reified *branch o*f and *tributary of* relationships in FMA and *isBranchOf* in GALEN. For example, from the fact that *Pharyngeal nerve* is a direct descendant of *Branch of maxillary nerve* is FMA, *<Pharyngeal nerve, branch of, Maxillary nerve>* was augmented in FMA. Since *branch of* (*tributary of, isBranchOf*) is not transitive, the augmentation is limited to direct descendants of *Branch of Y*.

o Identifying relationship patterns
Relationship patterns represent the correspondence between relationships (or combination thereof) across systems. Such patterns are identified by investigating the relationships among anchors in the two systems. More precisely, for each associative relationship between two anchors in one system, we searched for all shortest paths between the same two anchors in the other system. Both hierarchical and associative relationships are allowed in the paths. However, we ignored the paths where an associative relationship and its inverse are present because such paths are usually not indicative of an associative relation of interest between the two anchors. An associative relationship between two anchors in system and a combination of relationships between the same two anchors in the other system compose a path pair. Concepts are removed from the paths to create relationship patterns. Additionally, these patterns are simplified by representing several successive relationships of the same kind by only one relationship. These transformations generate pattern pairs from path pairs. For example, from the path pair:
FMA: *Pancreas → arterial supply → Dorsal pancreatic artery*
GALEN: *Pancreas → isServedBy → CaudalPancreaticArtery → isBranchOf → InferiorPancreaticArtery → isBranchOf → DorsalPancreaticArtery*
the pattern pair *{FMA: arterial supply, GALEN: isServedBy - isBranchOf}* is obtained. The frequency of each pattern pair (i.e., the number of paths pairs this pattern pair comes from) was recorded in order to select only the most frequent pairs (as they are also expected to be the most significant ones), thus ignoring "accidental" pattern pairs.

❑ Results of comparing associative relationships
847 inter-anchor associative relations were identified in FMA and 6,922 in GALEN. About 10% of them were generated by augmentation in FMA and nearly 1% in GALEN. 4,070 inter-anchor path pairs between FMA and GALEN were obtained. 350 pattern pairs were identified from these path pairs. In 4% of the cases, two associative relationships match across systems (e.g., *{FMA: nerve supply, GALEN: isServedBy}*). In 87% of the cases, one associative relationship corresponds to a combination of hierarchical and associative relationships in another system (e.g., *{FMA: arterial supply, GALEN: isServedBy - isa}*. In 9% of the cases, pattern pairs consist of one associative relationship in one ontology and a hierarchical relationship in the other (e.g., *{FMA: bounded by, GALEN: haspart}*). Not surprisingly, about 56% of the associative relationships in FMA and 84% in GALEN do not appear in the patterns (e.g., *isPositionedDistalT*o in GALEN and *fascicular architecture* in FMA). The pattern pair with the highest frequency (13% of 4,070 path pairs) is *{FMA: partof, GALEN: isBranchOf}*. The pattern pair with the second highest frequency (8%) is *{FMA: branch of, GALEN: isBranchOf}*.

❑ Problems revealed while investigating associative relationships
o The augmentation of associative relationships in FMA identified a cycle between *<Common palmar digital vein, isa, Tributary of superficial palmar venous arch>* (from which the augmented relation is *<Common palmar digital vein, tributary of, Superficial palmar venous arch>*) and *<Common palmar digital vein, tributary, Superficial palmar venous arch>*.
o Two inverse patterns in one system matching the same pattern in another system may be indicative of the differences in the definition of the concepts identified as anchors. Take the two pattern pairs for example: *{GALEN: isSpaceDefinedBy, FMA: partof}* and *{GALEN: isSpaceDefinedBy, FMA: haspart}*. Path pair example for the first pattern is *<OrbitalCavity, isSpaceDefinedBy, Orbit>* in GALEN and *<Orbital cavity, partof, Orbit>* in FMA. Path pair example for the second pattern is *<ConjunctivalSac, isSpaceDefinedBy, Conjunctiva>* in GALEN and *<Conjunctival sac, haspart, Conjunctiva>* in FMA. The difference came from the fact that in FMA *Conjunctival sac* was not taken as a cavity while it was in GALEN.

Investigating the implicit knowledge representation in the FMA and GALEN [3]

- ❑ Investigating the contribution of techniques for acquiring implicit knowledge.
  In the explicit and implicit knowledge for identifying structural evidence for anchors, relations come from different sources: explicit, complemented, augmented or inferred. Some relations can only come from one source, while others come from multiple sources. We differentiate five such sources, defined as follows.
  - ○ Base semantics (i.e., valid relations either explicit or complemented, with hierarchical cycles and duplicates removed), denoted as B.
  - ○ Augmented semantics, denoted as A.
  - ○ Inferred semantics based on the base, denoted as I_B.
  - ○ Inferred semantics based on the augmented, denoted as I_A.
  - ○ Inferred semantics based on the base or augmented, denoted as $I_{B \cup A}$.

  Every relation is specified by all of its possible sources, i.e., at least one and at most five. Sources specified are those contributing to the relation, and sources not specified are those that can not generate the relation. For example,

  <center>*<Finger, PART-OF, Free upper limb>* (B, A, $I_{B \cup A}$)</center>

  means that the relation is in the base semantics, can be augmented, can be generated by $I_{B \cup A}$, and can not be generated by I_B or I_A. By identifying the source of the relations, we investigated the speciality of the explicit representation and implicit knowledge generated by various techniques, within and across systems.

- ❑ Some results about the contribution of techniques for acquiring implicit knowledge.
  From the perspective of the semantic relations, the source of a relation represents the method (or methods) by which this relation can be generated. From the five individual methods we studied (B, A, I_B, I_A, and $I_{B \cup A}$), nineteen sources in FMA and sixteen in GALEN were found to partition the total set of relations into disjoint subsets. To each subset corresponds a combination of methods by which the relations in the subset can be generated. Four sources contribute the vast majority of relations in both FMA (about 95%) and GALEN (nearly 99%). These sources are: $(I_{B \cup A} \cap I_B)$, $(I_{B \cup A})$, (B), and $(B \cap I_{B \cup A} \cap I_B)$. The number and percentage of relations coming from each source for FMA and GALEN are analyzed. For example, 105,084 relations in FMA can be generated by both A (augmentation) and $I_{B \cup A}$ (inference based on the base and augmented relations), but not by the other three methods. Moreover, we went into the base semantic relations, augmented relations and inferred relations in FMA and GALEN to analyze the proportion of relations that can be generated from different techniques. For example, the base semantic relations come from all sources involving B, including, for example, $(B \cap I_{B \cup A})$. While some of these relations are only present in the base (about 54.9% in FMA and 86.5% in GALEN), some of them may also be augmentable (6.7% in FMA and 2.7% in GALEN), be inferable (38.8% in FMA and 11% in GALEN), or both (0.5% in FMA and 0.2% in GALEN).

Evaluating our alignment approach by comparing to schema matching techniques [4,7]

- ❑ Comparing to the schema matching algorithm for aligning FMA and GALEN
  Unrelated to our work, Peter Mork et al., at Microsoft also took up the challenge of aligning FMA and GALEN. PM used Microsoft's generic schema matching algorithm. The methods consist of lexical mapping, structural mapping and hierarchical mapping. Each phase generated a separate set of matches with similarity scores in (0,1] and improved the similarity scores of some of the matches from the former phases. No distinction was made between concepts and relationships. Although both used lexical and structural techniques in general, PM's approach is domain-independent (thus directly applicable to other domains), while ours takes advantage of the domain knowledge to maximize the chances of finding the intersection as well as conflicts between the two representations. We compared the matches obtained from our approach with those from PM's.
  - ○ Aligning FMA [version July 2, 2002] and GALEN [v.4]
    For the purpose of comparing the two alignments, we applied our methods to the same versions of FMA and GALEN PM used.

- The first step was to identify anchors (i.e., shared concepts across systems) lexically through exact match and after normalization. Both preferred concept names and synonyms were used.
- The second step was to identify additional anchors through UMLS synonyms. Two concepts across systems were considered a match if their names share the same CUI in UMLS under the Anatomy semantic group. Again, both preferred concept names and synonyms were used.
- The third step was to identify anchors structurally. Anchors received no, positive or negative structural evidence, respectively, if they share no, some, or conflicting hierarchical connections to other anchors. Both *IS-A* and *PART-OF* hierarchical relationships were used. Connections can be either direct or indirect.
- The fourth step was to compare the associative relationships across two systems based on the anchors supported by positive structural evidence.

o Comparing our alignment with Peter Mork's alignment
Our alignment is composed of pairs of concepts with structural evidence and pairs of associative relationships across systems. PM's alignment is composed of pairs of concepts/relationships with similarity scores across systems. PM considered as valid only the matches with similarity scores higher than or equal to some specified threshold. We compared these two alignments to get their intersection and the matches specific to each alignment. Moreover we analyzed the reasons that caused the shared and exclusive matches and suggested some possible improvements for each method.

❑ Results of aligning FMA [version July 2, 2002] and GALEN [v.4]
Our lexical alignment identified 2410 concept matches across systems. Through UMLS synonyms, 366 additional matches were obtained, resulting in totally 2776 matches. By structural alignment, 2536 (91.4%) out of 2776 matches received positive evidence, 40 (1.4%) negative evidence and 200 (7.2%) no evidence. PM identified 3780 matches. 2583 (68.3%) of them have similarity scores higher than or equal to the threshold set as 0.83, and 2539 out of 2583 actually have the similarity score 1.0. There are 126 matches related to relationships (*PART-OF* and associative). 102 of them are one relationship matching one concept, which constitute incorrect pairs, and 24 are relationship pairs.

❑ Results of comparing the two alignments
o 2448 concept matches are common to the two alignments.
- 36 of them received negative structural evidence by our approach but all have similarity scores higher than or equal to the threshold by PM's. This is because PM didn't identify conflicts between two systems.
- 171 received no structural evidence by ours; 168 of them have similarity scores higher than or equal to the threshold by PM's.
- 2241 received positive structural evidence by ours. 2199 of them have similarity scores higher than or equal to the threshold by PM's, so these are the really shared matches of the two alignments. 42 of 2241 have similarity scores lower than the threshold, partly because that PM didn't use synonyms of FMA in the lexical mapping.
o 328 concept matches are specific to our alignment
- 12 of them are matches from preferred terms.
- 168 are from synonyms in FMA, which were not used by PM.
- 148 were matched through UMLS synonyms. PM missed them because of the difference the two approaches have in using UMLS synonyms. In our approach, UMLS synonyms were compared against the whole concept name while partial matches were allowed by PM's approach.
o 1332 matches are specific to PM's alignment
- 17 of them are relationship pairs with similarity scores higher than or equal to the threshold. 3 of 17 are of *PART-OF* relationships, 7 were also identified by our associative relationship comparison and 7 not.
- 72 of them are of anonymous concepts in GALEN. So far our approach hasn't dealt with anonymous concepts yet.
- 1111 of them are non-anonymous concept pairs with similarity scores lower than the threshold.
- 132 of them are non-anonymous concept pairs with similarity scores higher than or equal to the threshold. Domain knowledge will be required to validate these matches.

Evaluating our alignment approach by aligning another two anatomical ontologies of different species [9]

We applied the techniques of identifying anchors lexically and structurally between FMA and GALEN to comparing another two representations of anatomy of different species. The first representation is the Adult Mouse Anatomical Dictionary (MA) describing the anatomical structure of the adult mouse. The second is the anatomy subset of NCI Thesaurus (NCI) for cancer research, describing the naturally occurring human biological structures, fluids and substances. Unrelated to our study, a manual alignment between MA and NCI was obtained by a domain expert at the Jackson Laboratory. We compared the matches obtained by our fully automatic approach and by manual approach. Moreover, we applied our structural alignment techniques to both groups of matches. Hierarchical relationships among matches were acquired in each system and compared across systems in order to identify positive structural evidence for the mapping. The results were reviewed by an expert. This work is a collaboration with the Jackson Laboratory and NCI.

Our fully automatic approach identified 715 matches based on lexical similarity, while the manual approach identified 781. 639 matches are common to both approaches. The vast majority of the 639 shared matches is supported by positive evidence, while only 45 (7%) of them received no evidence. After review by an expert, it was determined that all of the 639 shared matches are valid, including the 45 without structural evidence. Secondly, 76 matches are specific to our automatic approach, not identified by the manual approach. 80% of them are valid according to the expert review. Lastly, 142 matches were only identified by the manual approach, and 94% of them were validated.

Aligning multiple ontologies [8]

Mappings among multiple ontologies can be built pairwise. Alternatively, one ontology can be selected as a reference and all other ontologies only need to be mapped to this reference. We investigated the indirect alignment of two anatomical ontologies through a reference ontology, and to compare it to direct alignment between these two ontologies. The ontologies under investigation are the Adult Mouse Anatomical Dictionary (MA) and the anatomy subset of NCI Thesaurus (NCI). The FMA serves as reference ontology. To our knowledge, this is the first attempt to derive mappings automatically among anatomical ontologies from the alignment of these ontologies to a reference.

❑ Methods

Our study consists of three phases: 1) three direct alignments: MA-NCI, MA-FMA, and NCI-FMA; 2) an indirect alignment between MA and NCI through their direct alignments with the FMA; and 3) a comparison of the direct alignment MA-NCI to the indirect alignment obtained through the FMA.

o Direct alignment

For the three direct alignments, we applied the techniques for identifying anchors lexically and structurally developed for aligning FMA and GALEN. In the lexical alignment, both preferred names and synonyms in MA, NCI and FMA are used. Additionally, UMLS synonymy was used to identify additional matches. In the structural alignment, inference rules were applied to three systems, and specifically in FMA we made explicit the reified *PART-OF* relations. Structural alignment identifies the similar hierarchical relationships among lexical matches as well as structural conflicts. For example, the concepts *Forelimb* in MA and *Upper extremity* (synonym: *Forelimb*) in NCI were identified as a match through synonymy. Moreover, the match is supported by positive structural evidence because they exhibit similar relations to other matches in the two systems, including *Limb* (through *IS-A*), *Arm* and *Hand* (through *HAS-PART*).

o Indirect alignment

An indirect mapping was derived automatically between MA and NCI from the two direct alignments MA-FMA and NCI-FMA. For example, the direct alignment MA-FMA identified a match {MA: *Forelimb*, FMA: *Upper limb* (synonym: *Forelimb*)}, which is supported by positive evidence. The direct alignment NCI-FMA identified a match {NCI: *Upper extremity*, FMA: *Upper limb* (synonym: *Upper extremity*)}, also supported by positive evidence Therefore, the match {MA: *Forelimb*, NCI: *Upper extremity*} is derived automatically, through the FMA concept *Upper limb*, supported by positive structural evidence in both direct alignments.

o Comparison of the direct and indirect alignments

We compared the matches obtained by direct alignment MA-NCI and by indirect alignment through the FMA. The matches were classified into three groups: matches identified by both alignments; matches

specific to the direct alignment MA-NCI; and matches specific to the alignment through the FMA. The match {MA: *Forelimb*, NCI: *Upper extremity*} belongs to the first group.

❑ Results
  o The three direct alignments
    The alignment NCI-FMA identified the largest number of matches (2173), MA-NCI the smallest (715), and MA-FMA in the middle (1353). A very small number of conflicts was identified in the two direct alignments to FMA, while none were identified in the direct MA-NCI alignment. In the three direct alignments, a vast majority of the matches (> 90%) was supported by positive structural evidence. No evidence (positive or negative) was found for 5-9% of the matches in three direct alignments.
  o Indirect alignment
    703 matches between MA and NCI were automatically derived from the 1353 matches in the direct alignment MA-FMA and the 2173 matches in NCI-FMA. 649 of them (92%) received positive structural evidence in both direct alignments MA-FMA and NCI-FMA, 8 (1%) received negative evidence in one of the two direct alignments, and 46 (7%) received no evidence in at least one of the two direct alignments.
  o Comparison of direct and indirect alignments
    We compared the 715 matches obtained in the direct alignment MA-NCI to the 703 matches resulting from the indirect alignment through the FMA. 654 matches are common to both alignments, leaving 61 matches specific to the direct alignment and 49 specific to the indirect alignment through the FMA. Among the 654 shared matches, 583 (89%) received positive structural evidence in all three direct alignments, e.g., {MA: *Forelimb*, NCI: *Upper extremity*}. 65 (10%) of the 654 shared matches received no structural evidence in at least one of the three direct alignments. Although linked to other matches in MA (e.g., *PART-OF Cranium*), *Chondrocranium* has no hierarchical relations to any other matches in NCI and FMA. Therefore the matches of *Chondrocranium* received no evidence in any of the three direct alignments. At last, 6 (1%) of the 654 shared matches received negative evidence in one of the three direct alignments. For example, while a concept *Pericardial cavity* is present in the three ontologies, the corresponding match received negative evidence in the direct MA-FMA alignment (from in *HAS-PART* relationship to *Pericardium* in MA while in *PART-OF* relationship to *Pericardium* in FMA), no evidence in MA-NCI, and positive evidence in NCI-FMA.

Enforcing ontology modeling principles [6]
❑ Assessing and enforcing compliance with ontological modeling principles in FMA
  We investigated the degree to which FMA complies with ontological principles. This is a case study of assessing and enforcing compliance with modeling principles in large-scale biomedical ontologies. Fifteen principles were selected for this purpose, covering various aspects of ontology modeling.
  o Principles about hierarchical cycles, e.g., neither *IS-A* or *PART-OF* cycles are allowed.
  o Principles about classification, e.g., every non-leaf concept has at least two children; in every group of siblings, each concept has specific properties or relations to other concepts; every non-leaf concept is classified according to a single criterion.
  o Principles about incompatible relationships, e.g., for every pair of concepts *X* and *Y*, *X* and *Y* do not have both *IS-A* and *PART-OF* relationships.
  o Principles about dependence, e.g., concept *Subdivision of X* does not exist unless concept *X* exists; anatomical concept with name containing "wall" is in *PART-OF* relationship with some larger concept.
  o Principles about co-dependence of equivalent relations, e.g., the co-dependence between equivalent relations <*X, IS-A, Subdivision of Y*> and <*X, PART-OF, Y*> must be identified.
  o Principles about implicit relations (i.e., those implicitly embedded in concept names or combinations of relations), e.g., the implicit relations are consistent with explicit relations.

❑ Results of assessing and enforcing compliance with ontological modeling principles in FMA
  FMA is found mostly in compliance with the principles we tested. This is significant considering FMA has been developed in frame-based ontology authoring tool Protégé without built-in mechanisms for specifying or checking the consistency. The lack of automatic support for enforcing ontological principles in the

authoring tool is the main reason why the system fails to completely comply with the principles. Some of the results are listed as follows.

- o Hierarchical cycles include:
  - 32 *PART-OF* cycles (direct or indirect) for 20 concepts.
  - 25 IP cycles (*X - IS-A -…- IS-A - PART-OF -…- PART-OF →X)* for 15 concepts.
  - 12 PI cycles (*X - PART-OF -…- PART-OF - IS-A -…- IS-A →X)* for 5 concepts.
- o Out of 23,368 non-leaf concepts, 99% (23,111) have at least two children while 1% (257) have a single child. 48% of 23,111 concepts with more than one child have at least two children sharing identical relations to other concepts. 16,181 concepts (70% of 23,111) have children whose names exhibit laterality by the words "left" and "right". 95% of 16,181 concepts are classified only by laterality criterion, while 5% exhibit multiple classification criteria.
- o Incompatible relationships include:
  - 309 pairs of concepts have both *IS-A* and *PART-OF* relationships (direct or indirect).
  - 21 pairs of concepts have both *IS-A* (direct or indirect) and *BRANCH-OF* / *TRIBUTARY-OF* (direct) relationships.
  - 430 pairs of concepts have both *PART-OF* (direct or indirect) and *BRANCH-OF/TRIBUTARY-OF* (direct) relationships.
  - 123353 pairs of concepts have more than one type of *PART-OF* relationships (direct or indirect).
- o 1,980 terms are in the form of "*Subdivision of X*" or "*Organ component of X*". In 91% cases, X is present in the system, while in 9% cases, X is missing. 1,321 terms contain the word "wall" corresponding to 1,068 concepts. 64% of these concepts are in some *PART-OF* relation, while 36% do not have any explicit *PART-OF* relationships to other concepts.
- o Among 147,077 direct explicit *PART-OF* relations (*<X, PART-OF, Y>*), 4% are equivalent to some *IS-A* relations (e.g., *<X, IS-A, Subdivision of Y>*), and 30% equivalent to some combinations of relations (e.g., *<X, PART-OF, Z>* and (*<Z, IS-A, Y>*).
- o 459,704 relations are explicitly represented in the system while 11,779,914 are embedded implicitly in concept names and combinations of relations. The implicit relations are generally consistent with themselves and with the explicit relations except in only 41 cases where cycles were found.

OWL-izing the FMA [11]

Collaborating with Dr C. Golbreich from the University of Rennes, France, we attempted to convert FMA [version 1.1, 2004] from its frame-based representation in Protégé to the description logic language OWL DL. The objective of this study is to investigate description logics versus frame-based formalisms in supporting large-scale biomedical ontologies.

- ❑ Extracting a subset of FMA for conversion
  The first problem met in the conversion is the size of FMA, whose classes, slots, and slot values add to nearly 190,000 frames in Protégé. For efficiency reasons, we extracted a subset of FMA for conversion. In our previous study, we noticed that a large proportion of the leaf concepts in FMA only bear laterality distinctions compared to their parents. For example, *Arm* has only two children, *Left arm* and *Right arm*, and they are leaf concepts. The extraction was performed by hiding these laterality leaf concepts (e.g., *Left arm* and *Right arm*) and the slot values the other concepts have in connection with them.

- ❑ Converting FMA from frame-based structure Protégé to description logic language OWL DL
  Protégé is equipped with OWL Plug-in with several automatic OWL conversion functions. However, they failed to handle FMA correctly due to its size and complexity. Therefore, we designed conversion rules and coded the conversion by ourselves. These conversion rules specify how to define the properties, classes and individuals in OWL based on the frame descriptions of FMA in Protégé.

  - o Properties in OWL

    Slots in Protégé with the types Integer, Float, String, or Symbol with values of true and false are converted to datatype properties in OWL. Slots in Protégé with the type Symbol and allowed parents or allowed values are converted to object properties in OWL, and the union of allowed entities becomes the range of the property. Slots manually defined as annotation are converted to annotation properties

in OWL. Additionally, single slots in Protégé are specified as functional properties in OWL while multi slots are not. Moreover, when a slot has inverse slot in Protégé, if the inverse is itself, then it is specified as symmetric property in OWL; otherwise it is specified as the inverse property of its inverse in OWL. Finally, for every slot, the union of all the classes having the slot in Protégé becomes the domain of the corresponding property in OWL.

- o Classes in OWL

  Concepts in Protégé are converted to classes in OWL, and become SubClassOf of their parent classes. Every slot value of the class in Protégé is converted to a necessary condition of the class in OWL by existential restriction (i.e., someValuesFrom) or hasValue restriction, depending on whether the slot is object or datatype property. For necessary and sufficient condition, we manually designated slots. If a class has such slot, then all the values of the slot are converted to one necessary and sufficient condition for the class in OWL by intersection of someValuesFrom restrictions. Additionally, when a class has a slot specified with allowed parents in Protégé, the union of all the allowed entities becomes a necessary condition of the class in OWL by universal restriction (i.e., allValuesFrom).

- o Individuals in OWL

  Each one of the allowed values of slots with the type Symbol in Protégé is converted to an individual of :Thing in OWL.

- ❑ Results of converting FMA to OWL DL

  About 40,000 FMA concepts and their slot values were extracted for conversion, accounting for 57% of originally 70,000 FMA concepts. For this extracted subset of FMA, the OWL file generated by the conversion contains 117,000 frames, including 40,000 classes and 77,000 conditions (necessary, and necessary and sufficient) for these classes. Additionally, there are 155 properties and 85 individuals specified in the OWL file. We designated the slot *constitutional part* for necessary and sufficient condition, resulting in 570 classes defined by necessary and sufficient conditions.

  The NCI Thesaurus in OWL is probably the largest file in the Protégé OWL repository so far. It contains 53,000 frames, including 34,000 classes and 9,000 conditions of these classes. This shows that it is much less complex and smaller than the FMA OWL. Moreover, there are no necessary and sufficient conditions, nor hasValue, allValuesFrom, or class union specified in the NCI Thesaurus, while all these restrictions occur in the FMA OWL. The complexity and size made the FMA OWL a big challenge for description logic systems including developing environments (e.g., Protégé OWL Plug-in) and reasoners (e.g., Racer).

- ❑ Classifying the FMA OWL by description logic reasoner Racer

  By reasoning in Racer, inconsistencies in the original FMA knowledge representation were identified which generally would have been missed in the frame-based system. Meanwhile, problems in our conversion from Protégé to OWL were revealed. These differences show the benefits of description logic systems over frame-based systems. On the other hand, the current reasoning software has technical limitations and it was impossible to reason about the FMA OWL as a whole. Therefore, specifications causing reasoning to fail had to be removed, e.g., global axioms (i.e., a class' necessary conditions other than its necessary and sufficient conditions). Additionally, optimization techniques were explored to downsize and simplify the FMA OWL file. Moreover, various smaller OWL files were generated for testing properties in groups (e.g., all datatype properties) or one by one. Some of the classification results and optimization techniques are described in the following sections.

  - o Unsatisfiable classes from reasoning on Boolean datatype properties

    Two classes conflict when they have true and false values respectively for the same Boolean datatype property. A class becomes unsatisfiable if it is descendant of both such classes. For example, class *Zone of cell* is identified as unsatisfiable by Racer because it is specified as not having mass and meanwhile a descendant of class *Material physical anatomical entity* which has mass. Additionally, descendants of unsatisfiable classes become unsatisfiable, and so do those having *someValuesFrom* restriction to unsatisfiable classes. As a result, a total of 116 classes in FMA OWL are identified as unsatisfiable by Racer from reasoning on Boolean datatype properties.

  - o Equivalent classes from reasoning on necessary and sufficient conditions

    In the original FMA in Protégé, classes are supposed to be distinct from each other. However, some classes are considered equivalent by Racer from reasoning on necessary and sufficient conditions on property *constitutional part*. For example, class *Wall of biatrial part of heart* becomes equivalent with

class *Wall of biventricular part of heart* because they are siblings and they have the same constitutional parts, resulting in sharing the same definition in FMA OWL. There are other cases where a class and its direct superclass become equivalent when they have the same constitutional parts. All these equivalences are semantically incorrect, which revealed the problem of using one uniform property for defining necessary and sufficient conditions of classes.

- o Transitivity of properties

  Defining a property as transitive may reveal more problems in class definitions in FMA OWL. For example, class *Cell* is defined as an anatomical structure and has three other classes as constitutional parts. When *constitutional part* is specified as transitive, some descendants of anatomical structure acquire the same three classes as constitutional parts through transitivity reasoning, thus becoming direct subclasses of *Cell*, including *Prostate*, *Cardiac atrium* and *Breast*. This semantically incorrect classification can be avoided if the definition of *Cell* is strengthened by adding more relations into its necessary and sufficient condition such as cell is bounded by surface of cell.

- o Complementing inverse relations

  When properties are specified as inverse of each other and used in combination with *someValuesFrom* restriction, some optimization mechanism in Racer has to be switched off which may cause the reasoning to fail. Therefore, in order to test inverse properties, we removed the *inverseOf* specification in FMA OWL, automatically added the missing inverse relations, and converted them in OWL the same way explicit relations were converted. Complemented relations may change the classification in Racer. For example, class *Set of ureteric branches of testicular arteries* is complemented to be arterial supply of *Ureter*. *Set of ureteric branches of testicular arteries* is specified as direct subclass of *Anatomical set template* but after classification it has changed its superclass to *Anatomical structure*. This is because *Set of ureteric branches of testicular arteries* is not in the domain of property *arterial supply of* (which is descendant of *Anatomical structure*) in the original FMA, but is inferred to be so because of complementation. This indicates either that the original *arterial supply* relationship between *Ureter* and *Set of ureteric branches of testicular arteries* is wrong, or the domain of *arterial supply of* should be made broader.

- o Unsatisfiable classes from reasoning on domain and range of properties

  The domain and range of a property in OWL can be used to infer that the class having the property is subclass of the domain and the class as target of the property is subclass of the range. Unsatisfiability occurs when the use of a class in the domain or range of property conflicts with the classification of the class. For example, class *Surface of wrist* becomes unsatisfiable when it is specified as having *someValuesFrom* property *2D part* to class *Anatomic snuff box*. The range of *2D part* is *Non-material physical anatomical entity*, but *Anatomic snuff box* is a descendant of *Material physical anatomical entity*. *Non-material physical anatomical entity* being disjoint with *Material physical anatomical entity* eventually made the class *Surface of wrist* unsatisfiable.

- o Optimizing the domain of a property
  The domain of a property in FMA OWL is the disjunction of all classes whose specification in Protégé contain the introduction of the corresponding property. The number of such classes may be very large. For example, the property *arterial supply* has 4007 classes in its domain, and *location* 1618. This makes reasoning in Racer very inefficient or even impossible. A domain can be optimized if classes which are descendants of others are removed. As a result, 40 properties had their domains downsized, e.g., *arterial supply* has only 4 classes left in its domain and *location* only 2. After the optimization, the size of domains of properties did not cause problems in Racer anymore.

Analyzing the spatial knowledge representation in the FMA using Region Connection Calculus (RCC model)

- ❑ About RCC model
  RCC model describes the spatial relationships between regions: *DC* (Dis-Connection), *EC* (External Connection), *PO* (Partial Overlap), *TPP* (Tangential Proper Part), *TPPi* (inverse of *TPP*), *NTPP* (Non-Tangential Proper Part), *NTPPi* (inverse of *NTPP*) and *EQ* (Equality). They are jointly exhaustive and pairwise disjoint (JEPD), which means that any two regions stand in exactly one of the right relationships. There is a (8 * 8) composition table for the RCC-8 relationships. *R3* is the composition of *R1* and *R2* means for any three objects *a, b, c*, if *R1(a,b)* and *R2(b,c)* hold, then *R3(a,c)* must hold. For example, for any *a, b,*

and *c*, if *TPP(a,b)* and *DC(b,c)*, then *DC(a,c)* holds. In most of the cases the composition results in a disjunction of two or more relationships. For example, for any *a, b* and *c*, if *PO(a,b)* and *TPP(b,c)*, then *(PO(a,c) ∨ TPP(a,c) ∨ NTPP(a,c))* holds. In a few cases the composition leads to nothing. For example, for any *a, b,* and *c*, if *DC(a,b)* and *DC(b,c)*, then *a* and *c* can be in any of the eight RCC relationships, which basically tells nothing, because this is already implied in the JEPD feature of RCC model.

❑ Converting the partitive and associative relations and attributed slot values in the FMA to RCC relationships
  o Converting rules
    The FMA is not built based on the RCC model. Instead, spatial connections among anatomical entities must be derived from partitive and associative relationships, e.g., *constitutional part of*, *contained in*, *surrounded by*, *branch of*, *muscle origin*, and *adjacent to*. Moreover, some attributed slots in the FMA also represent spatial anatomical knowledge. Here is an example of attributed slot *location*:
        Heart|location|laterality(Right)|coordinate(Lateral)|adjacent(true)|related object(Right lung)|
    It represents the spatial connections between *Heart* and *Right lung*: they are adjacent. A manual analysis was performed by a domain expert, and resulting in more than 50 conversion rules. Some of these rules are shown in the following:

    constitutional part of->TPP ∨ NTPP

    constitutional part->TPPi ∨ NTPPi
    continuous with->EC
    branch of->EC
    adjacent to->EC
    muscle origin->EC

    contained in->PO ∨ TPP ∨ NTPP ∨ EQ

    surrounded by->EC ∨ EQ

    attributed regional part->TPPi ∨ NTPPi
    location + adjacent(true)->EC
    location + adjacent(false)->DC

  o Methods of converting
    Conversion rules were applied to partitive and associative relations and attributed slot values. For any two concepts *X* and *Y*, if they don't have any relation that can be converted into RCC relationships, then *RCCcnvt(X,Y) = ALLRCCS =DC ∨EC ∨PO ∨TPP ∨TPPi ∨NTPP ∨NTPPi ∨EQ*. If *X* and *Y* has only one relation *<X, R, Y>* that can be converted, and slot *R* has conversion rule $R->rc_i ∨ … rc_n$, then we have *RCCcnvt(X,Y) = rc_i ∨ … rc_n.* If *X* and *Y* stand in another relationship *R'*, and $R'->rc_j ∨ … rc_m$, then *RCCcnvt(X,Y) = (rc_i ∨ … rc_n) ∧ (rc_j ∨ … rc_m)*. Then we check if *RCCcnvt(X,Y) = ⊥*, which indicates some inconsistency in the direct FMA spatial knowledge representation between *X* and *Y*, or inaccuracy in the conversion rules of slots involved.

  o Results of converting
    From conversion, 84284 pairs of concepts have obtained RCC relationships (excluding *ALLRCCS*), of which 18112 have only one RCC relationship. 1818 pairs of concepts are in ⊥, and the causes for inconsistency need to be further investigated.

❑ Applying the RCC composition table to infer additional RCC relationships among concepts in the FMA
  o Methods of applying RCC composition table
    For any two concepts *X* and *Z,* if there does not exist a concept *Y* such that neither *RCCcnvt(X,Y)* nor *RCCcnvt(Y,Z)* is *ALLRCCS* or ⊥, then *RCCcnvt(X,Z) = ALLRCCS*. If there exists such a concept *Y*, then according to the RCC composition table, we get *RCCcomp(X,Z) = CT(RCCcnvt(X,Y), RCCcnvt(Y,Z))*. If there exists another such class *Y'*, then *RCCcomp(X,Z) = CT(RCCcnvt(X,Y), RCCcnvt(Y,Z)) ∧ CT(RCCcnvt(X,Y'), RCCcnvt(Y',Z))*. We check if *RCCcomp(X,Z) = ⊥*, which indicates some inconsistency in the FMA spatial knowledge representation between *X* and *Z* indirectly through other classes, or inaccuracy in the conversion rules of slots involved.

Note: $CT((rc_k \vee \ldots rc_s), (rc_t \vee \ldots rc_l)) = \vee ct(rc_i, rc_j)$, where rc_i loops from rc_k to rc_s and rc_j loops from rc_t to rc_l, and $ct(rc_i, rc_j)$ represents the composition result of two RCC relationships rc_i and rc_j.

Finally, for any two concepts *X* and *Y*, their final RCC relationships are $RCC(X,Y) = RCCcnvt(X,Y) \wedge RCCcomp(X,Y)$. We check if $RCC(X,Y) = \perp$. If so, and $RCCcnvt(X,Y) \neq \perp$ and $RCCcomp(X,Y) \neq \perp$, then this indicates an inconsistency between direct spatial knowledge of X and Y and their inferred spatial knowledge through other concepts, or inaccuracy in the conversion rules of slots involved.

- ○ Results of applying RCC composition table

  From composition, 707284 pairs of concepts have obtained RCC relationships (excluding *ALLRCCS*). 280822 pairs of concepts have obtained ALLRCCS and they were ignored. 698588 pairs out of 707284 acquired RCC relationships specifically from conversion, while 8696 pairs also acquired RCC relationships from conversion.

  For 698588 pairs obtaining RCC relationships only from composition, 28042 have only one RCC relationship. 1520 pairs are in $\perp$, and the causes for inconsistency need to be further investigated.

  For 8696 pairs obtaining RCC relationships both from conversion and composition, 1126 of them have only one RCC relationship in the final RCC. 2194 pairs are in $\perp$ finally and the causes for inconsistency need to be further investigated.

- ❑ Analysis of positive results

  Concept pairs with only one RCC relationship reveal definite spatial relationships between the two concepts. When two concepts do not have any direct slots that can be converted into RCC, if there exists a third concept that has RCC relationships with these two concepts respectively, then by composition we can infer the RCC relationships between these two concepts. This demonstrates the power of spatial reasoning by using RCC model. For example, *Esophagogastric junction* and *T10 part of heart* do not have any slots can be converted. However, *Esophagogastric junction* stands in *regional part of* relation to *Abdominal part of esophagus*, which is converted into proper parts in RCC. On the other hand, *T10 part of heart* has attributed slot *location* to *Abdominal part of esophagus* with "*adjacent(false)*", thus it was inferred that the two concepts are disconnected in RCC. The reasoning is shown as follows.

  ∵ *RCCcnvt(Esophagogastric junction, Abdominal part of esophagus) = TPP $\vee$ NTPP*
    RCCcnvt(Abdominal part of esophagus, T10 part of heart) = DC

  ∴ *RCC(Esophagogastric junction, T10 part of heart) = CT(RCCcnvt(Esophagogastric junction, Abdominal part of esophagus), RCCcnvt(Abdominal part of esophagus, T10 part of heart))*
  = CT((TPP $\vee$ NTPP), DC) = DC

- ❑ Analysis of negative results

  Concept pairs with $\perp$ in their RCC relationship indicate either inconsistencies in the FMA spatial knowledge representation or inaccuracies in our conversion rules. For example, we found many cases where two concepts are continuous with each other and, simultaneously, have a common third concept as part. For example:

  RCCcvnt(Wall of right side of heart, Fibroelastic connective tissue of endocardium) = TPPi $\vee$ NTPPi, from relationship *constitutional part*

  RCCcvnt(Wall of left side of heart, Fibroelastic connective tissue of endocardium) = TPPi $\vee$ NTPPi, from relationship *constitutional part*

  RCCcnvt(Wall of right side of heart, Wall of left side of heart) = EC, from relationship *continuous with*

  Therefore, through concept *Wall of left side of heart*, one conjunct in *RCCcomp(Wall of right side of heart, Fibroelastic connective tissue of endocardium)* is *CT(RCCcnvt(Wall of right side of heart, Wall of left side of heart), RCCcvnt(Wall of left side of heart, Fibroelastic connective tissue of endocardium)) = CT(EC, (TPPi $\vee$ NTPPi)) = DC $\vee$ EC*. Since *(DC $\vee$ EC) $\wedge$ RCCcvnt(Wall of right side of heart, Fibroelastic connective tissue of endocardium) = $\perp$*, the final *RCC(Wall of right side of heart, Fibroelastic connective tissue of endocardium) = $\perp$*.

Publications:

1. Songmao Zhang, Olivier Bodenreider, Knowledge augmentation for aligning ontologies: An evaluation in the biomedical domain. *Proceedings of the Semantic Integration Workshop at the Second International Semantic Web Conference (ISWC 2003)* 2003:109-114
2. Songmao Zhang, Olivier Bodenreider, Aligning representations of anatomy using lexical and structural methods. *Proceedings of AMIA Annual Symposium* 2003:753-757
3. Songmao Zhang, Olivier Bodenreider, Investigating implicit knowledge in ontologies with application to the anatomical domain. In: Altman RB, Dunker AK, Hunter L, Jung TA, Klein TE, editors. *Pacific Symposium on Biocomputing 2004*: World Scientific; 2004:250-261
4. Songmao Zhang, Peter Mork, Olivier Bodenreider, Lessons learned from aligning two representations of anatomy. In: Hahn U, Schulz S, Cornet R, editors. *Proceedings of the First International Workshop on Formal Biomedical Knowledge Representation (KR-MED 2004)*; 2004:102-108
5. Songmao Zhang, Olivier Bodenreider, Comparing associative relationships among equivalent concepts across ontologies. *Medinfo 2004*:459-463
6. Songmao Zhang, Olivier Bodenreider, Law and Order: Assessing and enforcing compliance with ontological modeling principles. *Computers in Biology and Medicine*; 2005:(accepted)
7. Songmao Zhang, Peter Mork, Olivier Bodenreider, Philip A. Bernstein, Comparing two approaches for aligning representations of anatomy. *Artificial Intelligence in Medicine*; 2005:(submitted)
8. Songmao Zhang, Olivier Bodenreider, Alignment of multiple ontologies of anatomy: Deriving indirect mappings from direct mappings to a reference. *AMIA*; 2005:(submitted)
9. Olivier Bodenreider, Terry Hayamizu, Martin Ringwald, Sherri De Coronado, Songmao Zhang, Of mice and men: aligning mouse and human anatomies. *AMIA*; 2005:(submitted)
10. Songmao Zhang, Olivier Bodenreider, Structural alignment of anatomical ontologies. 2005:(in preparation)
11. Songmao Zhang, Olivier Bodenreider, Frame-based structures vs. description logics in representation and reasoning of the Foundational Model of Anatomy. 2005:(in preparation)

Presentations:

- A framework for comparing medical knowledge representation systems. Mid-term presentation as participant in the Medical Informatics Training Program at the NLM, December 11, 2002
- Aligning representations of anatomy using lexical and structural methods. Talk for 2003 summer rotation students at the NLM, June 25, 2003
- Aligning representations of anatomy using lexical and structural methods. Paper presentation at *AMIA 2003*, Wasgington DC, November 13, 2003
- Investigating implicit knowledge in ontologies with application to the anatomical domain. Paper presentation at *Pacific Symposium on Biocomputing 2004*, Hawaii, January 6, 2004
- Lessons Learned from aligning two representations of anatomy (presented by Peter Mork). Paper presentation at *KR-MED 2004*, Whistler, Canada, June 1, 2004
- Comparing associative relationships among equivalent concepts across ontologies. Paper presentation at *Medinfo 2004*, San Francisco, September 11, 2004
- Aligning representations of anatomical knowledge. Final presentation as participant in the Medical Informatics Training Program at the NLM, March 30, 2005
- The body region connection calculus (presented by Olivier Bodenreider). Workshop on mapping the human body, University of Buffalo, April 17, 2005